Historical and architectural context for traffic management needs today

kc claffy
The Cooperative Association for Internet Data Analysis
The University of California, San Diego
kc@caida.org

January 27, 2010

Reply Comments NBP Public Notice # 30.
These comments refer to GN Docket Nos. 09-47, 09-51, and 09-137.

Abstract

In December I was asked to present for a technical advisory panel at the FCC on the topic of Internet traffic management [10]. This document expands on some points I made there, in response to questions from the FCC during the panel, as well as other panels in December and January. I begin with historical context for the constraints on our traffic management concepts and capabilities today. Recognizing the reality that someone needs to pay for infrastructure, I emphasize the importance of measurable transparency in protecting private property rights as well as individual (consumer) rights, including requiring objective data to demonstrate the need for traffic management approaches that might restrict the freedom of users. Two international case studies (Japan and Canada) illustrate how others are confronting the same regulatory questions; both examples suggest that the FCC is headed in the right direction regarding transparency obligations, including requiring measurement tools to enable consumer awareness of their own traffic patterns. The FCC is also fostering more enlightened regulatory capabilities in the future by establishing a longer-term advisory function to provide empirically grounded analysis of the (predicted) success or failure of (proposed) policies.

1 My background

I have been studying various Internet research topics since 1990. In 1993 I co-authored my first paper on proposed traffic management approaches to deal with congestion, in an interdisciplinary paper entitled, "Mitigating the Coming Internet Crunch", in collaboration with the National Science Foundation and a management professor [22]. In 1994 I published my doctoral dissertation on "Internet traffic characterization" [5], using public traffic data whose collection was mandated by the U.S. government. (A thesis that cannot be reproduced today, unless you're in Japan, where researchers have come closest [15].) I recently wrote "Ten Things Lawyers Should Know about Internet Research" [6], which covers these and other issues related to broadband policy. In May 2009 I presented highlights of this piece to the FCC [16].

2 Historical perspective

A few minutes of review of historical facts about the Internet architecture provide illuminating context for the pace and limitations of core architectural innovation, including support for traffic management. In 1966, Larry Roberts published a paper, "Towards a Cooperative Network of Time-Shared Computers" [23], which led to DOD commissioning the construction of the ARPANET in 1969. The ARPANET, in turn, became the technical core of the current Internet. In 1977, Leonard Kleinrock from UCLA published a paper titled "Hierarchical Routing for large networks; performance evaluation and optimization" [17], the techniques in which bear strong similarity to core routing technology still used on the Internet today.

In 1980, ARPANET had its first system-wide failure, grinding to a complete halt due to a "statusmsg" virus, not due to an attack but rather a configuration error. Undeterred, and unable to reign in growth and interest in connectivity to ARPANET, another U.S. federal agency, the National Science Foundation, agreed in 1986 to build a larger more general-purpose, higher performance network that could connect a much broader R&E community than ARPANET's mission permitted. NSF also launched a program to fund the middle mile – the "NSF-funded regionals" – geographically limited networks that handled connecting up individual campuses to the core NSFNET backbone. The same year, the Internet Engineering Task Force was established. To keep up with growing demand, in 1991 the NSFNET backbone upgraded to 45Mbps, 30 times the current bandwidth, and officially allowed commercial institutions to connect to academic institutions via the NSFNET backbone. A commercial interconnection facility emerged for use in connecting emerging commercial Internet networks to others. Even at this time there were issues with traffic management of applications not originally envisioned on the NSFNET backbone, in particular real-time (streaming) audio and video across the NSFNET, which prompted our paper on how to mitigate an expected imminent bandwidth crunch [22] (see section 5).

By 1995 there was substantial commercial network activity, which the U.S. government did not see as appropriate to compete with, so NSF orchestrated a stable transition of the users the NSFNET backbone, including the regional networks, from the NSFNET to private sector Internet access service providers. Not even a decade passed before *The Economist* posted a cover story: "How the Internet killed the phone business" [3]. By this time the incumbent local access monopolies in the U.S. had mostly bought or developed their own Internet access business, swallowing the growing independent competition. In an ironic twist of economic fate, a subset of the Bell Operating Companies created by the 1984 breakup of AT&T repurchased the divested AT&T in 2005, reconstituting much of the former Bell System. Notably, throughout this decade, the incumbent telecom industry lobbyists convinced the U.S. courts and regulators to abandon many common carriage and essential facilities obligations that had always been part of U.S. communications policy, over and above the "regulatory forbearance" philosophy espoused in the 1996 Telecommunications Act.

3 What didn't change?

Several fundamental aspects of the Internet architecture have not changed over the last two decades of astonishingly dynamic history.

1. The Internet still uses a *network architecture*, i.e., the Internet Protocol (IP) suite, built for a cooperative file-sharing environment and relatively low bandwidth applications [28]. Not only are there no mechanisms for communicating quality of service requirements across multiple different service providers, but there are also no means (nor incentive) for enforcement of such requirements, which would be essential to scalable commercial deployment.

2. We are still using a *routing architecture* that makes certain assumptions about the characteristics of the underlying network topology in order to optimize efficiency and performance. Specifically, Kleinrock's

1977 technique [17] of aggregating nearby nodes into groups, these groups into larger groups, etc., is the basis of popular hierarchical approaches used today for both interdomain and intradomain routing. Unfortunately, these assumptions about hierarchical structure in the topology have become less true as the Internet has evolved [18], leaving us with a looming architectural problem which no proposed alternative has yet solved.

3. Because it is an integral part of the network architecture, we are still using the same IP addressing architecture, despite extensive and elaborate (and thus far failed) efforts to upgrade to an addressing architecture (IPv6) with sufficient address space to meet expected global needs in the 21st century [7].

4. We are still using roughly the same *transport architecture* [29], although experimentation increases, motivated by the need to efficiently and fairly support P2P protocols, e.g., the LEDBAT effort [2].

5. The host *naming architecture* has not fundamentally changed since the DNS was introduced in 1983, although significant modifications have attempted to accommodate new functionality such as IPv6 [19], and secure naming [1].

6. Several fundamental aspects of the economic architecture also seem invariant, including that after two decades we do not seem to have a sustainable competitive business model for transmitting bits across long distances in an increasingly ubiquitously connected world. More to the point, the privatized platform for bit delivery still exhibits natural monopoly economics, even 13 years after the U.S. legislated otherwise.

7. Wire (and wireless) spectrum allocations remain determined by the same tiny handful of facilities owners with monopoly power over network access.

8. Our ability to make progress – or even quantitatively assess – the "4 S's" of critical infrastructure: security, scalability, sustainability, stewardship, is characterized by painfully incremental progress, stunted by lack of transparency into the infrastructure.

4 What did change?

Over the same four decades, remarkable changes did happen while the core Internet architecture stabilized (to some "ossified" [24]) for two reasons: the need for reliability as the infrastructure began to subsume all other communication media, and economic forces that impeded architectural disruption (to some "innovaton" [24]) across many competing entities.

- *Industry structural trajectory and capital reserves.* The Internet has exhibitied a stark contrast to telephony in its political history. Telephony started as a private sector invention, but once the (U.S.) government recognized it as potentially critical general purpose infrastructure, they imposed heavy regulation to ensure broad (later universal) accessibility and other socially desirable functionality, e.g., 911. The Internet followed the opposite historical timeline: during its first 30 years it was almost completely funded, developed, managed, and operated by federal government agencies and their awardees, but once the government recognized it as potentially critical general purpose infrastructure, they removed existing and avoided new regulation as much as possible.

- *Bandwidth provisioning efficiency*, which exhibited an exponential increase annually for at least a decade, the product of fantastic advances in optical multiplexing technologies.

- *Data processing/storage efficiency* benefited from even more rapid technological advancement [27].

- *Access provisioning and peering models* transitioned from being relatively transparent under government-supported infrastructure, to private, unregulated, and opaque, unamenable to objective empirical macroscopic analysis, in parallel with it becoming critical infrastructure.

- *Naming provisioning* (DNS registration) privatized with the rest of the infrastructure in the mid-1990s, and is now supported by a competitive industry subject to lightweight – and some argue ineffective, from a security perspective – regulation by ICANN.

- *Address provisioning*, originally handled by U.S. government contract to a single administrative entity, transitioned to a participatory, transparent, needs-based governance regime (the Regional Internet Registries, or RIRs) in the mid-1990s. As the RIRs have confronted their own unfortunate failure to steward a transition to IPv6, several RIRs are now launching a new regime of private ownership of IPv4 addresses, some members still hoping that IPv6 will happen eventually.

- *Pricing models* are difficult to analyze at the wholesale level due to their treatment as trade secrets (see peering, above), and at the retail level pricing has monotonically increased (in the U.S.) as competition has decreased. No surprise there. More surprising are the public admissions by commercial providers that bit transport must *"compete with other forms of telecommunications.. including things such as DVD distribution via the mail" [11]*, without acknowledging that the U.S. Postal Service is bound by a public charter that dictates profit minimization, while carriers are bound to profit maximization.
- *Data access*: freely available from the NSFNET backbone in the 80s and 90s, today limited data is released, to a select few researchers for specific purpose, under strict NDA
- *Innovative uses of the network* have only just begun.

5 What did we recommend in 1994?

Our 1994 paper offered a simple, cooperative solution, using existing fields of the IP packet, and some wildly academic assumptions about incentives to cooperate with eachother and respect them. The full abstract was:

> *The current architecture and implementation of the Internet assumes a vast aggregation of traffic from many sources and stochastic distribution of traffic both in space (traffic source) and time (burstiness of traffic volume). Given this general assumption, Internet components typically have little if any ability to control the volume and distribution of incoming traffic. As a result the network, particularly from the perspective of the router, is vulnerable to significant consumption of networking resources by high-volume applications, with possibly little stochastic behavior, from a few users. This often impacts the overall profile of network traffic as aggregated from many clients. An example is the continuous flows introduced by real time applications such as packet audio, video, or rapidly changing graphics.*
>
> *This situation creates a time window where applications exist on a network not designed for them, but before an appropriately*

architected network can augment the current infrastructure and cope with the new type of workload. We propose a scheme for voluntarily setting Internet traffic priorities by end-users and applications, using the existing 3-bit Precedence field in the Internet Protocol header.

Our proposal has three elements. First, network routers would queue incoming packets by IP Precedence value instead of the customary single-threaded FIFO. Second, users and their applications would voluntarily use different and appropriate precedence values in their outgoing transmissions according to some defined criteria. Third, network service providers may monitor the precedence levels of traffic entering their network, and use some mechanism such as a quota system to discourage users from setting high precedence values on all their traffic. All three elements can be implemented gradually and selectively across the Internet infrastructure, providing a smooth transition path from the present system. The experience we gain from an implementation will furthermore provide a valuable knowledge base from which to develop sound accounting and billing mechanisms and policies in the future.

Even in 1994, we did warn that then current architecture was living on borrowed time, and that our naive academic proposal would only be a temporary measure until new architectural capabilities were developed. Neither our nor any subsequent cooperative inter-domain (i.e., competitive) solutions ever got traction; their most lasting effect was to convince operators that academics were out of touch with industry reality, especially the economics. It is worth noting that for intra-domain traffic, i.e., that owned and operated by a single administrative domain, the "reasonable traffic management" technology problem was solved and deployed long ago – across multiple domains, the problem is not the technology.

The question of inter-domain quality of service (QoS) for traffic management came up at December's technical advisory panel [10], when Walter Johnston asked about the discrepancy between the IETF and ITU positions on QoS. He cited a debate from a summit several years ago where the biggest departure between the ITU and IETF perspectives was on QoS, with the IETF arguing that the Internet was all about "best effort", and the ITU counter-arguing that interdomain QoS was essential. I responded that they were both right from each perspective – the IETF was speaking from an network architecture perspective, while the ITU was speaking from a "how to stay in business" perspective, at least insofar as how they defined their business. Note that the IETF, despite cultural loyalty to their best-effort philosophy, spent a decade developing standards for interdomain-QoS technical solutions, while the ITU worked on its proprietary IMS solution behind closed doors. But neither group has yet succeeded in providing a scalable, sustainable solution for interdomain QoS. That is to say, none of this technology ever worked in the marketplace, for entirely technical reasons, including Dave Clark's frequently quoted explanation, "we never learned how to route money."

6 What have we learned about QoS technology and economics?

The academic research community is in a rather absurd situation of not being able to do even the most basic network research, even on the networks established explicitly to support academic network research. This limitation has led to unresolvable contradictions in our field, including on the most politically relevant network research question of the decade: what are the costs and benefits of using QoS to support multiple service classes, to users as well as providers, and how should these service classes be determined? Two research papers that contradict eachother on this topic illustrate the problem. Internet2 has stated that QoS will never make sense to deploy on its backbone [25], apparently based on the (unpublished) economics of Internet2's infrastructure. In contrast, AT&T and collaborators have argued that QoS is critical, although they have offered no data to support or validate their claims [14].

Scientific network researchers have not solved the "empirical validation" problem for much of their discipline. Several funding agencies have realized the depth of the problem, and its impact on national security and public safety, and tried to address it via technical means, e.g., supporting research on data anonymization. But since the problem is more about policy and incentives than technology, these efforts have had limited impact. More recently, DHS has launched an effort to articulate a set of ethical principles and guidelines for network research [12], ultimately seeking a community-ratified framework that will help advance privacy-respecting yet empirically grounded Internet research.

7 What conditions did we (and then NTIA) recommend in 2009?

On 23 March 2009 I spoke at an NTIA "Roundtable on Nondiscrimination and Interconnection Obligations" [9], where I was asked to comment on what conditions should accompany the Broadband Stimulus awards, in particular what data obligations should award recipients have. I listed the most data-deprived, i.e., opaque, areas of interdomain Internet ecosystem dynamics: penetration, peering, performance, and pricing. With respect to tiered pricing as a part of traffic management policies, I agreed with Andrew Odlyzko's cogent assessment, *"To evaluate claims about need for additional revenues...one needs solid cost data and a dynamic model of the industry. At the moment we do not have either one." [20]*

I also listed the most data-deprived dimensions of critical infrastructure conditions: security, scability, sustainability, and stewardship. Progress in all of these areas will resist evaluation, must less assistance, without additional transparency into the evolution of economics, traffic, topology, and routing. Given limited FCC/NTIA resources and capabilities, I emphasized the need to leverage other sources of capital as well as intellectual resources.

The language on data collection that survived the sausage-making process into the Notice of Funding Availability (NOFA) for the $7B broadband funds, was:

> *Awardees receiving Last Mile or Middle Mile Broadband Infrastructure grants must report, for each specic BTOP project, on the following:*
>
> *i) The terms of any interconnection agreements entered into during the reporting period;*
>
> *ii) Traffic exchange relationships (e.g., peering) and terms;*
>
> *iii) Broadband equipment purchases;*
>
> *iv) Total & peak utilization of access links;*
>
> *v) Total & peak utilization on interconnection links to other networks;*
>
> *vi) IP address utilization & IPv6 implementation;*
>
> *vii) Any changes or updates to network management practices.*

At the December FCC panel [10], Jon Peha asked specifically what data was needed to effectively make a decision on traffic management regulation.

I responded that NTIA's list would go a long way, but as in the Japan and Canada case studies presented (sections 9 and 10), the burden of data provision should be on the provider to demonstrate the need for traffic management. The operators know how they run their networks, and what they want to do to manage them, so the first step is to have them propose data that would prove their case, and let the FCC's newly formed independent technical advisory panel help evaluate whether they have sufficiently proved it. Recently two European researchers (Wagter and Felten) issued a challenge to ISPs to provide a dataset to allow the researchers to (pro bono) analyze the role and impact of heavy users on bandwidth consumption and performance. They offered a detailed specification of a dataset they could usefully analyze [26]. I suspect no ISP will volunteer (i.e., invest in people time to prepare the specified) data.

8 In lieu of data... Prevailing risks

Regulators should be aware of several red flags associated with allowing tiered or metered pricing of Internet bandwidth. First, it is important to recognize that Internet links are typically filled with a substantial fraction of pollution, or garbage traffic unwanted by the recipient. Spam is only one such category of traffic, but much more unwanted traffic crosses the Internet, even destined for addresses that have no machines assigned to them, much less email accounts. Figure 1 shows a sample of traffic to a large chunk of empty IPv4 address space, i.e., addresses with no machines attached to them. The chart shows several gigabytes per hour are being sent (from all across the global Internet) to this chunk of unused addresses. Other studies of traffic going to empty space show similar levels of traffic, although with different characteristics [21]. This prevailing condition on the Internet sheds some doubt on pricing plans based on metering traffic, as much of it is not caused – or desired – by the end consumer. Regulators should certainly not be allowing metering or caps such as that announced by Comcast in 2008, with no tool for users to measure and analyze their own bandwidth consumption. (Comcast released a trial usage measurement tool to Portland customers in December 2009, but no details or feedback have been available [8] to evaluate it.) In addition, there should be substantial disclosure to consumers of how their Internet service is affected by packet prioritization. (Comcast made an initial attempt at such disclosure [8] following the public outcry and FCC criticism in 2008.)

More fundamentally, regulators must be careful not to incentivize busi-

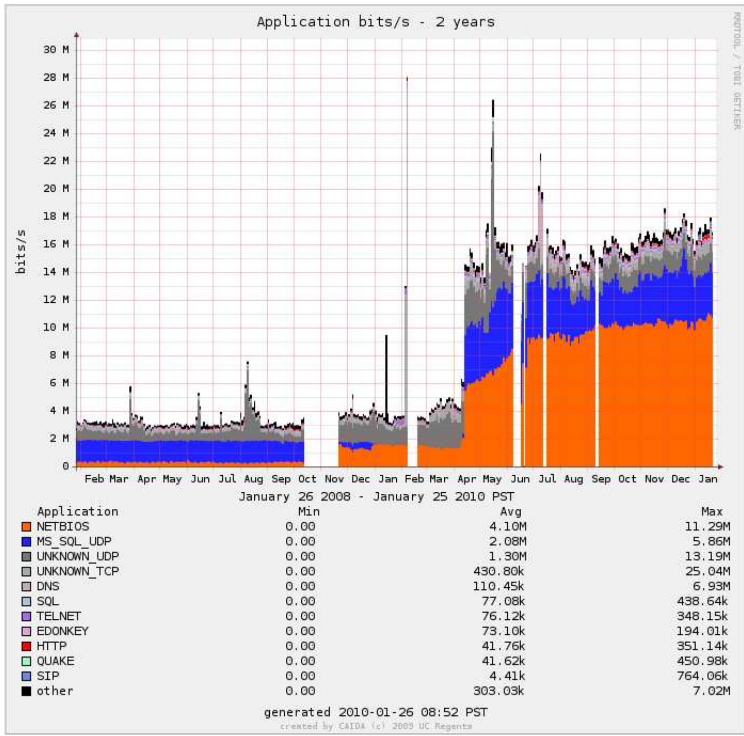

Figure 1: *Bits per second to a large chunk (/8) of empty address space, 22008-2010*

ness strategies based on technologies to induce artificial scarcity, as opposed to pro-growth & innovation-driven strategies. Traffic metering and price tiering can have an inherently warping effect on provisioning incentives, since consumers who are buying a 20Mbps service for \$20/month are unlikely to ever increase their tier unless they are throttled, perhaps artificially, at the 20Mbps rate. The irresistible temptation (if not fiduciary obligation to Wall Street) to throttle under these circumstances may induce an exploitation scenario, most pronounced when the sum of tiered prices extracted could easily cover the capex cost of expanding the maximum aggregate bandwidth rate, but the upgrade does not happen, so consumers are just stuck with inferior service indefinitely. Avoiding this socially inferior outcome requires limiting "reasonable" bandwidth reservations & tiering scenarios to situations where (a) the bandwidth tiers sum to approximately the upper bandwidth limit of current technology, *and* (b) the sum of the billing combined does not exceed the actual cost of service delivery by multiple orders of magnitude.

9 Japan Case Study

It is also worth examining international case studies from countries who have recently confronted the traffic management issue. In May 2008 a group of Japanese industry associations collaborated on set of "Guideline for Packet Shaping" [13], to propose rough industry consensus on circumstances that render packet shaping acceptable. Analyzing specific illustrative examples, the document does an impressive job of clarifying the relationship between "secrecy of communications"[1] and fairness in use under Japanse business law. They agreed that packet shaping should be implemented only in exceptional circumstances, either to "facilitate necessary network management" or "protect users". In operational terms, they agreed that traffic shaping must fit three criteria: (1) it must be in response to congestion of specific heavy users that is degrading, or is likely to degrade service of general users; (2) must be substantiated by objective data; and (3) it must use a method that aims only at the necessary objective, i.e., controlling congestion. They admit the terminology is necessarily vague, so they provide case studies, including one where they illustrate why content examination, e.g., looking for copyright infringement based on payload, is *not deemed reasonable*, because one cannot do it accurately for, nor scalably for all users. They also point out that other traffic shaping measures, e.g, protection against inappropriate P2P software behavior, should be explicitly about protecting users, and require informed consent of the user.

The Japanese report acknowledges several areas that need additional research to enable enlightened policy: impact of increased video content; impact of packet shaping on access network bandwidth; application-specific packet-shaping; paid peering for content, which they acknowledge as problematic as ISP's themselves expand into content; information-sharing among players regarding packet shaping implementation; and P2P protocol efficiency improvement.

10 Canada: Case Study

The Canadian Radio-television and Telecommunications Commission released a set of guidelines in October 2009 that requires Internet service

[1]Japanese law has a strong (broader than in U.S. privacy law) notion of "secrecy of communications", which can include content, names, locations, timestamps, headers, and other artifacts of individuals. Acts of "infringement" include intentionally gaining knowledge of matters that fall under secrecy of communications and using the knowledge to one's own or another's interests against the parties of original communication.

providers to be more transparent about their Internet traffic management practices [4]. Their conclusions and recommendations are not only similar to the spirit of FCC's reactions and positions on this topic thus far, but congruent with Japan's position outlined above. In particular, they require traffic management practices to: (1) minimize harm to user (so, don't block when delaying will work); (2) be based on a transparent need (must show data); and (3) be narrowly tailored (technically efficient) to accomplish the traffic management objective and not beyond. They acknowledge that the challenge will ultimately be about defining "reasonable" practice, but emphasize two policy principles: (1) targeting specific content or applications is not allowed; and (2) techniques must be based on not just quantifiable, but actually quantified, data.

References

[1] `http://www.ietf.org/dyn/wg/charter/dnsext-charter.html`.

[2] Ietf low extra delay background transport (ledbat) ietf working group. `http://www.ietf.org/dyn/wg/charter/ledbat-charter.html`.

[3] How the Internet killed the phone business. *The Economist*, 2005. `http://www.economist.com/displaystory.cfm?story_id=4400704`.

[4] Canadian Radio-television and Telecommunications Commission. Telecom Regulatory Policy CRTC 2009-657: Review of the Internet traffic management practices of Internet service providers, October 2009. `http://www.crtc.gc.ca/eng/archive/2009/2009-657.htm`.

[5] Kimberly Claffy. *Internet Traffic Characterization*. PhD thesis, 1994. Computer Science and Engineering Department, `http://www.caida.org/publications/papers/1994/itc/`.

[6] Kimberly Claffy. Ten Things Lawyers should know about Internet research, August 2008. `http://www.caida.org/publications/papers/2008/`.

[7] Kimberly Claffy. Report from the first workshop on internet economics (wie2009). *Computer Communications Review*, 2010. `http://www.caida.org/publications/papers/2010/wie_report/`.

[8] Comcast. Comcast Network Management Policy, 2009. `http://networkmanagement.comcast.net/`.

[9] Federal Communications Commission. Broadband Technology Opportunities Program (BTOP), Public Meeting Nondiscrimination and Interconnection Obligations, March 2009. `http://www.ntia.doc.gov/broadbandgrants/meetings.html`.

[10] Federal Communications Commission. Workshop: Technical Advisory Process Workshop on Broadband Network Management, December 2009. `http://www.openinternet.gov/workshops/technical-advisory-process-workshop-on-broadband-network-management.html`.

[11] Cogent Communications. Cogent Communications Group Q3 2009 Earnings Call Transcript, October 2009. http://seekingalpha.com/article/172306-cogent-communications-group-q3-2009-earnings-call-transcript?page=-1.

[12] Erin Kenneally and Michael Bailey and Doug Maughan. A Framework for Understanding and Applying Ethical Principles in Network and Security Research. In *Workshop on Ethics in Computer Security Research (WECSR 2010)*, 2010. http://www.caida.org/publications/papers/2010/framework_ethical_research/.

[13] Japan Internet Providers Association (JAIPA), Telecommunications Carriers Association (TCA), Telecom Services Association (TELESA), and Japan Cable and Telecommunications Association (JCTA). Guideline for Packet Shaping, May 2009. http://www.jaipa.or.jp/other/bandwidth/guidelines_e.pdf.

[14] Joseph D. Houle and K. K. Ramakrishnan and Rita Sadhvani and Murat Yuksel and Shiv Kalyanaraman. The Evolving Internet - Traffic, Engineering, and Roles. 2008. http://web.si.umich.edu/tprc/papers/2007/786/Evolving Internet.pdf.

[15] Hiroshi Esaki Kenjiro Cho, Kensuke Fukuda and Akira Kato. The Impact and Implications of the Growth in Residential User-to-User Traffic. *ACM SIGCOMM*, 2006. http://www.iijlab.net/k̃jc/papers/rbb-sigcomm2006.pdf.

[16] Kimberly Claffy. Ten Things Lawyers should know about Internet research, May 2009. http://www.caida.org/publications/presentations/2009/top_ten_fcc/.

[17] L. Kleinrock and F. Kamoun. Hierarchical routing for large networks: Performance evaluation and optimization. *Computer Networks*, 1:155–174, 1977.

[18] D. Krioukov, kc claffy, K. Fall, and A. Brady. On compact routing for the Internet. *Comput Commun Rev*, 37(3):41–52, 2007.

[19] M. Crawford and C. Huitema. DNS Extensions to Support IPv6 Address Aggregation and Renumbering, July 2000. http://www.ietf.org/rfc/rfc2874.txt.

[20] Andrew Odlyzko. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. Technical report, School of Mathematics and Digital Technology Center, University of Minnesota, 2008.

[21] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet Background Radiation. In *Internet Measurement Conference (IMC)*, 2004.

[22] R. Bohn and H.-W. Braun and K. Claffy and S. Wolff. Mitigating the coming Internet crunch. Technical report, UC, San Diego and National Science Foundation, 1994.

[23] Larry Roberts and Thomas Marill. Toward a Cooperative Network of Time-Shared Computers. October 1966. http://portal.acm.org/citation.cfm?id=1464336.

[24] Steve Bellovin, David Clark, Adrian Perrig, Dawn Song. A Clean-Slate Design for the Next-Generation Secure Internet, July 2005. NSF workshop report, http://www.nsf.gov/cise/cns/geni/ngsi.pdf.

[25] Ben Teitelbaum and Stanislav Shalunov. Why premium ip service has not deployed (and probably never will). Technical report, Internet2 QoS Working Group, 2006. http://qos.internet2.edu/wg/documents-informational/20020503-premium-problems-non-architectural.html.

[26] Herman Wagter and Benot Felten. Dataset Specification for Disruptive Broadband User Analysis, 2009. http://harmonica.typepad.com/Dataset_Specification_for_Disruptive_Broadband_User_Analysis_v1.pdf.

[27] Chip Walter. Kryder's law. *Scientific American*, August 2005. http://www.scientificamerican.com/article.cfm?id=kryders-law.

[28] Wikipedia, 2009. http://en.wikipedia.org/wiki/Internet_Protocol_Suite.

[29] Wikipedia, 2009. http://en.wikipedia.org/wiki/Transmission_Control_Protocol.